

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

ALAN V.

OK. In the last several lectures, we've been discussing the Fourier and Z-transforms. And

OPPENHEIM:

we've seen in particular that the Fourier and Z-transforms provide us with a set of important analytical tools for representing discrete time signals, and also for dealing with discrete time systems. We saw, for example, that through the use of the Fourier transform or the Z-transform, we could convert convolution in the time domain to multiplication in either the frequency domain in the Fourier transform case, or more generally, in the Z domain in the Z-transform case.

Now one of the things that it's important to recognize is that for the most part, the Fourier transform and the Z-transform are primarily analytical tools. That is, it would be hard to imagine implementing, for example, a discrete time system by first computing the Fourier transform of the sequences, multiplying the Fourier transforms together, and then computing the inverse transform.

One of the reasons that that's obviously a difficult thing to do computationally is that we saw, as we discussed, for example, the Fourier transform, that the Fourier transform is a function of a continuous variable. That is, ω in the Fourier transform is a continuous variable. So that, in fact, if we wanted to compute the Fourier transform explicitly, we would have to compute it at an infinite number of frequencies. Similarly, we have a situation like that for the Z-transform.

Well, today I'd like to introduce a third transform, which I'll refer to as the discrete Fourier transform. The discrete Fourier transform is similar in style to the Fourier transform and the Z-transform, as we've been talking about, in the sense that more or less, the discrete Fourier transform maps convolution to multiplication. Also, there are properties of Fourier transform-- the discrete Fourier transform-- that are similar to the properties of the Fourier transform and Z-transform, as we've been talking about them.

But the discrete Fourier transform is different than the other transforms that we've been discussing in a number of important respects. One of the respects, one of the reasons that it's different is that it is a transform that can be explicitly evaluated. And consequently, it is important, not only in the analysis of discrete time systems, but also, as we'll see, in the implementation of discrete time systems. That is, many digital signal processing algorithms, as

we'll see, actually involve the explicit computation of the discrete Fourier transform, which is the Fourier transform that we're about to introduce.

Now let me try to explain a little about what the discrete Fourier transform is before we look at it in detail. Basically, the discrete Fourier transform is related to the Fourier transform that we've been discussing in the sense that it corresponds to samples of the Fourier transform, or more generally, samples of the Z-transform. Now that requires that we impose some restrictions on the sequences that we're representing through that transform. And as we'll see, it turns out that we can represent sequences that are of finite length. That is, only have a finite number of non-zero samples. We can represent those sequences by samples of their Fourier transform. Those samples then correspond to the discrete Fourier transform.

Well, there are some issues that arise in looking at the discrete Fourier transform, or DFT, as I'll refer to it. And a number of these issues relate to the fact that the discrete Fourier transform has properties that are somewhat different, and also somewhat similar, to the Fourier transform properties that we've been discussing. There are lots of ways of introducing the discrete Fourier transform. And the one that I guess I find the most interesting, and the most satisfactory, is to relate the discrete Fourier transform to the discrete Fourier series for periodic sequences. And in particular, relate the notion of finite length sequences to periodic sequences.

Well, let me explain in a little more detail what I mean by that. Let's consider, first of all, a sequence x of n , which I'll restrict to be a finite length sequence. The sequence, for example, one example indicated here, has a set of non-zero values only over a finite range of the argument n . Here is a sequence that is of finite length. And I've chosen to refer to it as a sequence of finite length capital N . 0 outside the range from little n equals 0 to little n equals capital N minus 1. It's 0 outside that range.

It's interesting, just as an aside, to note that for this particular sequence, it's also actually 0 outside the range from little n equals 1 to capital N minus 1. And in general, obviously, if I talk about a finite length sequence of length capital N , I can also refer to it as a finite length sequence of length greater than capital N . That is, the important statement about the sequence being finite length, of finite length capital N , is that the sequence values are 0 outside the range 0 to capital N minus 1. Although obviously, the sequence could also be 0 inside that range for some of the values.

Now the basic notion that leads to the discrete Fourier transform, or one way of looking at the discrete Fourier transform, is to recognize that if I have a finite length sequence, as I have here, I could construct from that sequence a periodic sequence. And let me denote the periodic sequence by \tilde{x} of n . In general, by the way, when I refer to a sequence with a tilde on it, that will always correspond to a periodic sequence.

And let me construct this periodic sequence by simply taking the finite length sequence and repeating it over and over again with a period of capital N . In other words, I can construct the periodic sequence, \tilde{x} of n equal to x of n plus x of n shifted to the left by capital N and x of n shifted to the right by capital N and $2n$, 2 capital N , and 3 capital N , et cetera. In other words, just simply taking this sequence and repeating it over and over again with a period of capital N .

So obviously I can generate, from a finite length sequence, a periodic sequence. And in fact, I could get the finite length sequence back from the periodic sequence by simply extracting one period of this periodic sequence. In other words, I can get x of n back from \tilde{x} of n simply by multiplying by unity for n between 0 and capital N minus 1 , and 0 outside that range.

Now the important point here-- there are a couple of important points. One of them is that if I have a finite length sequence, I can turn it into a periodic sequence. If I have a periodic sequence, I can get back to the finite length sequence simply by extracting one period. Or what that essentially says is that there really isn't much difference between a finite length sequence and a periodic sequence. That is, a finite length sequence is defined by capital N values.

A periodic sequence is also defined by capital N values because once I specify a single period, then I don't have any more degrees of freedom in specifying the rest of the sequence. And that's an important point to keep in mind, particularly as we go through the next several lectures. A finite length sequence is very similar to a periodic sequence in that both of them are simply defined by capital N values. One way to think of that, by the way, is to think of taking a finite length sequence of capital N values. Instead of drawing it along a straight line as I've done here, imagine taking this finite length sequence and wrapping it around the circumference of a cylinder.

So we start with the finite length sequence and just simply display it wrapped around the circumference of a cylinder. As we run around the cylinder, over and over again, what we see is the periodic sequence \tilde{x} of n . So in a sense, the periodic sequence is just simply like the

finite length sequence, but wrapped on a cylinder instead of laid out in a straight line. And that's also a picture that will recur several times as we go through the next several lectures. So we can generate the periodic sequence from the finite length sequence.

We can also recover the sequence x of n from the periodic sequence by extracting just one period, as I've indicated here, x of n is \tilde{x} of n in the range, little n between 0 and capital N minus 1. And it's equal to 0, otherwise. Or what we'll use as a convenient form of notation for that is to express x of n , the finite length sequence, as the periodic sequence, \tilde{x} of n multiplied by another sequence, which I'll denote as R sub capital N of n , where R sub capital N of n is just simply a rectangular sequence.

So R sub capital N of n is 1 for little n between 0 and capital N minus 1, and it's equal to 0, otherwise. All right. So it's important to begin, at this point, to think of finite length and periodic sequences as more or less the same type of thing in the sense that it's easy to go back and forth from one to the other.

Now, why is this point of view important? Well, we know certainly in the continuous time case that a periodic time function can be represented by a Fourier series. In the discrete time case, a periodic sequence, likewise, can be represented by a Fourier series. And the idea, that is, the key point behind the discrete Fourier transform is that we can use the Fourier series representation of the periodic sequence to represent the finite length sequence. That is, that, in essence, provides a Fourier kind of representation for a finite length sequence.

So we have the notion then that the periodic sequence \tilde{x} of n has a Fourier series representation. We can compute the discrete Fourier series of this periodic sequence. And as we'll see-- not in this lecture, but as we'll see in more detail in the next lecture-- it is discrete Fourier series representation of this periodic sequence that is what we'll call the discrete Fourier transform of x of n .

So let's begin then with a discussion of the discrete Fourier series of periodic sequences with the idea that we'll be applying that representation to the representation of finite length sequence. And that representation is what will correspond to the discrete Fourier transform. OK. So we want to talk about the discrete Fourier series. We're considering a sequence \tilde{x} of n , which is periodic, and it's period we'll call capital N .

In the continuous time case, we know that if we have a periodic time function, we can represent it as a linear combination of harmonically related complex exponentials. That's the

Fourier series representation in the continuous time case. And I'll just simply state without proof that the same kind of relationship holds in the discrete time case. That is, we can represent a periodic sequence $x[n]$ as a linear combination of complex exponentials, which are harmonically related to the frequency.

Or equivalently, the reciprocal of the period, so that this forms a general relationship for a discrete Fourier series of a periodic sequence $x[n]$. That is, these are the harmonically related complex exponentials, just as we form linear combinations of harmonically related complex exponentials for the Fourier series in a continuous time case. What are the Fourier coefficients? Well, it's, of course, these capital $X[k]$, which we'll have a little more to say about in a minute.

Well, notice that I haven't specified as of yet any limits on this summation. And in particular, what we need to examine is how many distinct periodically or harmonically related complex exponentials there are. Well, let's take a look at the complex exponential, the set of complex exponentials $e^{j\frac{2\pi}{N}nk}$. And the statement that I want to make is that these complex exponentials are periodic. Of course, we know that they're periodic in n . But they're also periodic in k .

As we vary k from 0 to capital N minus 1, we generate all of the possible harmonically related complex exponentials with this fundamental frequency, $\frac{2\pi}{N}$. Well, we can see that very simply by substituting in for k , $k + N$. And then recognizing that we can break this complex exponential into the product of two complex exponentials $e^{j\frac{2\pi}{N}nk}$ and $e^{j\frac{2\pi}{N}n(Nk)}$.

Well, these capital N 's cancel each other out. This factor is then $e^{j2\pi n}$. Well, any integer multiple of 2π up here then just simply reduces this to unity. So that, in fact, $e^{j\frac{2\pi}{N}n(k+N)}$ is the same as $e^{j\frac{2\pi}{N}nk}$. Well, that shouldn't be surprising, actually, because that's a point that's come up time and again as we've been going through these lectures.

The point being that sinusoids in the discrete time domain, as we vary sinusoids in frequency-- we've seen time and again, in the range 0 to 2π -- in fact, those are all the sinusoids that we can generate. And if we keep going in frequency, we just simply see the same ones over again. So this is a consequence of that.

But for the discrete Fourier series, it says an important thing. It says that in forming the

discrete Fourier series, once we've used the complex exponentials for k between 0 and capital N minus 1, we've used all the complex exponentials with this fundamental frequency that we have. And if we keep going with k , we're just simply going to see the same complex exponentials over and over again.

What does that say about the discrete Fourier series? It says that in the representation of the discrete Fourier series, the limits on this sum don't range from minus infinity to plus infinity. They run simply from 0 to capital N minus 1. So once I've looked at these linear combination of these complex exponentials for k between 0 and capital N minus 1, there are no new complex exponentials with that fundamental frequency that I'm able to find.

So this then is the form of the discrete Fourier series. There's one additional insertion that I'd like to make. And this is just simply a normalization factor. I want to multiply this by 1 over capital N . Obviously, that doesn't make any essential difference. It's just a factor for normalization. And it plays a role which is similar in the continuous time case to the role that 2π or 1 over 2π plays.

All right. So here we have the discrete Fourier series. It's relatively straightforward to show that you can obtain the Fourier series coefficients \tilde{x} of k from \tilde{x} of n through an inverse relationship, which is the relationship that I've indicated here. So this is then, in essence, the inverse discrete Fourier series, or equivalently, the relationship for obtaining the discrete Fourier series coefficients from the periodic sequence \tilde{x} of n .

Now notice that I've happened to write the Fourier series coefficients with a tilde over them, implying that those coefficients are themselves periodic. Well, are they periodic, or aren't they periodic? What I've been saying here and what I've been saying here-- and I've spent a long time saying this-- that there are only a finite number of complex exponentials. Once I've looked at them in the range 0 to capital N minus 1, I've seen all the ones I can see.

And in that sense, the Fourier series coefficients, \tilde{x} of k , are finite. That is, there are only a finite number of them. Well, the important point is that there are only a finite number of distinct coefficients. Also in this relationship, whether I consider these as periodic or not periodic, I still only use a finite number of them. That is, I only use them in the range k equals 0 to capital N minus 1. And how I choose to interpret capital \tilde{x} of k outside the range 0 to capital N minus 1 is going to have absolutely no effect on the evaluation of this summation.

Well, it turns out to be convenient to interpret the Fourier series coefficients as being periodic in k . Well, in fact, the relationship as I've written it here makes it evident, in this particular relationship, that the coefficients are periodic. In other words, if I substitute in for k , k plus capital N , then I'll get back exactly the same relation that I have here. Because of the fact that e to the minus $j 2 \pi$ over capital N times little n times k plus capital N is equal to e the minus $j 2 \pi$ over capital N times little n times k .

So that if I simply examined-- I asked from this relationship-- what does capital X tilde of k plus capital N come out to be? If I simply substitute that in, then because of the fact that these two complex exponentials are the same, I'll find that x -- capital X tilde of k is equal to capital X tilde of k plus capital N , k plus 2 capital N , et cetera. That is, it's a periodic sequence, although I only use one period of it in reconstructing the periodic sequence x tilde of n .

I choose to do that mainly because of duality. That is, mainly because it is convenient to think of these as a periodic sequence so that I have one periodic sequence representing another periodic sequence, this being a periodic sequence in k with a period of capital N , and this being a periodic sequence in n with a period also of capital N . So now with a discrete Fourier series, there's a duality in the, what we could call the time domain and the frequency domain.

And the duality is there in part because we've chosen to represent, or think of the Fourier series coefficients as periodic, as a periodic sequence, although we only use a finite number of those values in actually explicitly evaluating the Fourier series for x tilde of n . Well, this then is the Fourier series. Let me finally rewrite it one other way, which just introduces some notation that's convenient.

Let me define w sub capital N as e to the minus $j 2 \pi$ over capital N . In that case, then just simply rewriting the Fourier series, we have x tilde of n is 1 over capital N , the sum of capital X tilde of k , w sub capital N to the minus nk . And the Fourier series coefficients expressed in terms of capital W sub n is the sum of x tilde of n , w sub capital N to the nk .

Well, the discrete Fourier series has properties, just as the Fourier transform and the Z -transform has had a number of properties. And again, as we've done with the other transforms, I won't spend a lot of time on the details of either enumerating the properties or proving them. But let me just illustrate one or two to give you some idea as to what these properties involve.

Well, first of all, as we've talked about in the Fourier transform and Z -transform cases, there is

a shifting property that tells us how the Fourier series coefficients of a periodic sequence are related to the Fourier series coefficients of that sequence shifted. And in particular, it turns out that if X_k and \tilde{X}_k are the Fourier series coefficients for $x[n]$, then the Fourier series coefficients for that sequence shifted, that is, $x[n+m]$, corresponds to multiplying the Fourier series coefficients by w_N^{-km} .

Shifting the sequence involves multiplying the Fourier series coefficients by a complex exponential, which is what this is. And that is similar to what we've seen with the Fourier transform, shifting a sequence, multiply the Fourier transform by a complex exponential. And the Z-transform, we had exactly the same situation. We have a dual relationship, which expresses the result of shifting the Fourier series coefficients. If we shift the Fourier series coefficients, the result is multiplication by a complex exponential of the original sequence, $x[n]$.

In fact, one of the things that's true with the discrete Fourier series that hasn't been true with the Fourier transform or the Z-transform is there is a strong duality between the time domain, the discrete time domain, and the Fourier or frequency domain. In particular, we begin with a discrete periodic sequence, and we end up in the Fourier domain with, again, a discrete periodic sequence.

In fact, something to think about is the fact that the Fourier series coefficients we've said, or we've chosen to interpret them, as being periodic. That implies that they themselves have a Fourier series representation. And something to think about is, what are the Fourier series coefficients of that periodic sequence?

All right. So one important point then is that we have this duality between the two domains. And essentially, any property that we have in the time domain will find the dual property in the frequency domain for the discrete Fourier series. And that will hold true when, later on in the next lecture, we talk about the discrete Fourier transform.

Another useful property, which we've also talked about for the Fourier transform, are the set of symmetry properties. And in particular, if we consider $x[n]$ to be a real a periodic sequence, then there are symmetry relationships between the real part of the Fourier series, and symmetry relations for the imaginary part of the Fourier series.

In particular, we can think of expressing the Fourier series coefficients in terms of their real part plus j times the imaginary part. And the symmetry relationships that result are that if the

periodic sequence is real, then the real part of the Fourier series coefficients are even, $x_{\text{R}}[k]$ is equal to $x_{\text{R}}[-k]$. And that's what we refer to as the property of, in the case of the Fourier transform, the Fourier transform being even.

We can also write this in another way that will become important when we talk about the discrete Fourier transform. In particular, since this is a periodic sequence, since $x[k]$ is periodic, obviously it's real part is also periodic-- we can replace k by k plus capital N , or k minus capital N . And we can rewrite this as a statement that says that $x_{\text{R}}[k]$ is equal to $x_{\text{R}}[N - k]$.

It shouldn't be particularly evident why we want to do that here, although it becomes important when we apply some of these notions to the representation of finite length sequence and sequences in the discrete Fourier transform. Likewise, for the imaginary part of the Fourier series coefficients, we end up with a symmetry property that says that the imaginary part of the Fourier series coefficients are odd. And again, as we did with the real part of the Fourier series coefficients, we can rewrite this to say that $x_{\text{I}}[k]$ is equal to $-x_{\text{I}}[N - k]$.

Finally, we can, in a similar manner, look at the magnitude of the Fourier series coefficients and the angle of the Fourier series coefficients. And just as we've seen with the Fourier transform, the result that we'll find is that the magnitude of the Fourier series coefficients are even, an even function of k . And the angle of the Fourier series coefficients are an odd function of k . So we have symmetry properties like we had with the Fourier transform.

An important thing to keep track of at this point, because we'll want to refer back to this when we talk about the discrete Fourier transform, is that in talking about the property of even or odd, we can either talk about it as we do here, or in terms of a shift-- that is, in terms of a statement that $x_{\text{R}}[k]$ is $x_{\text{R}}[N - k]$, which essentially relates the evenness of this periodic sequence to the relationship between values within one period.

Finally, we have another very important property, which is the convolution property. This is a property that, again, we've had for the Fourier transform and for the Z-transform. It's a property that states, essentially, that the convolution of two periodic sequences results in the multiplication of the discrete Fourier series coefficients, with just one minor twist, which again will become important when we relate this to the discrete Fourier transform.

In particular, we have a periodic sequence, x_1 tilde of n . And here its Fourier series coefficients. A second periodic sequence, x_2 tilde of n , with its Fourier series coefficients. And what we'd like to ask is, what is the periodic sequence, x_3 tilde of n , whose Fourier series coefficients are the product of x_1 tilde of k and x_2 tilde of k ? In other words, if we multiply the Fourier series coefficients together, what is the sequence that that corresponds to? The answer, which involves just simply a little bit of algebra to verify, is that the resulting sequence is a sum of x_1 tilde of m , x_2 tilde of n minus m .

Well, that looks exactly like a convolution with one important difference. And that is that the limits on the summation don't go from minus infinity to plus infinity as they did in the case of the Fourier transform and Z-transform. Ordinary linear convolution, as we've always been talking about it, involved the sum from minus infinity to plus infinity, of x_1 of m , x_2 of n minus m . Here we have, as the limits on the sum, 0 to capital N minus 1. The summation is carried out only over one period.

We have also a dual property, in other words, the property that relates the Fourier series coefficients of the product of two sequences to the Fourier series coefficients of the individual sequences. And because of the duality now that we have between the time domain and the frequency domain for the Fourier series, these Fourier series coefficients are, again, the convolution of the two Fourier series coefficients for x_1 of n and x_2 of n . There's a normalization factor, 1 over capital N . But again, the limits on the sum involve 0 to capital N minus 1.

So this is slightly different than the convolution as we've been talking about. It is referred to generally, and we'll be referring to it as a periodic convolution. The important distinction being that for this convolution, it involves a summation, not for minus infinity to plus infinity, but simply a summation over only one period of the periodic sequences. Well, to drive this home, let me just show you with a simple example what these sequences look like, and in particular, what the shifting involves, and apply an interpretation to that. And so let's return to the view graph.

What I'm indicating here is a sequence x_2 tilde of m , a sequence x_1 tilde of m , and to form the convolution of these two sequences. What I would like to do is generate x_2 tilde, in general, of n minus m , multiply that by this, by x_1 tilde of m , and then sum up the result over one period. Well, I'm indicating, by the way, in blue, just one period of this periodic sequence. And similarly in blue, one period of this periodic sequence.

So let's examine, first of all, the sequence which corresponds to replacing this by x^{2n-m} . And let's do it for the specific case where little n is equal to 2. So I've illustrated here the sequence x^{2n-m} . And to generate this sequence from this one involves essentially two steps. The first step is to flip this sequence over, that's replacing m by $-m$. And then the second step is to shift the sequence by an amount little n , depending on the argument that we want to stick in here.

Now this is the result of doing that. We can think of this as having flipped this sequence over, and then shifted it by two points to the right. And the result is then this periodic sequence. An important thing to keep in mind, or to look at-- and again, this is a point that we'll be emphasizing in much more detail in the next lecture-- is that as we examine these blue points, we could think of having gotten those by simply flipping one period of this and circularly shifting that, circularly shifting those points, simply in the range 0 to $N-1$. That's an interpretation that will become much more evident in the next lecture.

But then to form the convolution of this sequence with this sequence, we then, after having constructed x^{2n-m} , carry out the multiplication of these two sequences. The result of that multiplication, I've indicated here, so that for this particular example, these three points get multiplied by these three. These four points, rather, get multiplied by these four values. The rest of the points in that one period get multiplied by 0. And then finally, to evaluate the result of the periodic convolution for little n equal to 2, we sum up those values in the range 0 to $N-1$.

So it operates very much like a linear or ordinary convolution, as we've been talking about over the last several lectures. The important difference is that the summation is carried out just simply over one period. All right. Well, that completes the discussion as we want to present it, of the discrete Fourier series. As I indicated at the beginning of the lecture, our objective was eventually to develop a Fourier representation for finite length sequences, that is, the discrete Fourier transform.

And in the next lecture, what we'll want to do is take the ideas of the discrete Fourier series, as we've talked about them here, and apply them to the representation of finite length sequences, resulting in what we'll call the discrete Fourier transform.